

ABSTRACT

In a database system, a method of maintaining a self-tuning histogram having a plurality of existing rectangular shaped buckets arranged in a hierarchical manner and defined by at least two bucket boundaries, a bucket volume, and a bucket frequency. At least one new bucket is created in response to a query on the database. Each new bucket is contained within at least one existing bucket and the new bucket becomes a child bucket and the existing bucket containing it becomes a parent bucket. The boundaries of each new bucket correspond to a region of the database accessed by the query and the frequency of the new bucket is a number of data records returned by the query. Buckets may be merged based on a merge criterion such as similar bucket density when the total number of buckets exceeds the predetermined budget. The boundaries of a new bucket may be shrunk if the boundaries of the new bucket intersect any existing bucket boundaries.